arXiv:1703.04143v1 [cs.GT] 12 Mar 2017

Bernoulli Factories and Black-Box Reductions in Mechanism Design *

Shaddin Dughmi[†], Jason D. Hartline[‡], Robert Kleinberg[§] and Rad Niazadeh[§]

[†]Department of Computer Science, University of Southern California,
[‡]Department of Computer Science, Northwestern University,
[§]Department of Computer Science, Cornell University.

March 14, 2017

Abstract

We provide a polynomial time reduction from Bayesian incentive compatible mechanism design to Bayesian algorithm design for welfare maximization problems. Unlike prior results, our reduction achieves exact incentive compatibility for problems with multi-dimensional and continuous type spaces.

The key technical barrier preventing exact incentive compatibility in prior black-box reductions is that repairing violations of incentive constraints requires understanding the distribution of the mechanism's output, which is typically #P-hard to compute. Reductions that instead estimate the output distribution by sampling inevitably suffer from sampling error, which typically precludes exact incentive compatibility.

We overcome this barrier by employing and generalizing the computational model in the literature on *Bernoulli Factories*. In a Bernoulli factory problem, one is given a function mapping the bias of an "input coin" to that of an "output coin", and the challenge is to efficiently simulate the output coin given only sample access to the input coin. Consider a generalization which we call the *expectations from samples* computational model, in which a problem instance is specified by a function mapping the expected values of a set of input distributions to a distribution over outcomes. The challenge is to give a polynomial time algorithm that exactly samples from the distribution over outcomes given only sample access to the input distributions.

In this model, we give a polynomial time algorithm for the function given by *exponential weights*: expected values of the input distributions correspond to the weights of alternatives and we wish to select an alternative with probability proportional to an exponential function of its weight. This algorithm is the key ingredient in designing an incentive compatible mechanism for bipartite matching, which can be used to make the approximately incentive compatible reduction of Hartline et al. (2015) exactly incentive compatible.

*Shaddin Dughmi was supported in part by NSF CAREER Award CCF-1350900. Jason Hartline was supported in part by NSF grant CCF-0830773. Robert Kleinberg was partially supported by NSF grant CCF-1512964. Portions of this work were completed while Dr. Kleinberg was a researcher at Microsoft Research New England. Rad Niazadeh was supported by NSF grant CCF-1512964 and by a Google Ph.D. Fellowship. Portions of this work were completed while he was an intern at Microsoft Research Redmond.

1 Introduction

We resolve a five-year-old open question from Hartline et al. (2011, 2015): *There is a polynomial time reduction from Bayesian incentive compatible mechanism design to Bayesian algorithm design for welfare maximization problems.*[1] The key distinction between our result and those of Hartline et al. (2011, 2015) is that both (a) the agents' preferences can be multi-dimensional and from a continuous space (rather than single-dimensional or from a discrete space), and (b) the resulting mechanism is exactly Bayesian incentive compatible (rather than approximately Bayesian incentive compatible).

A mechanism solicits preferences from agents, i.e., how much each agent prefers each outcome, and then chooses an outcome. *Incentive compatibility* of a mechanism requires that, though agents could misreport their preferences, it is not in any agent's best interest to do so. A quintessential research problem at the intersection of mechanism deign and approximation algorithms is to identify black-box reductions from approximation mechanism design to approximation algorithm design. The key algorithmic property that makes a mechanism incentive compatible is that, from any individual agent's perspective, it must be *maximal-in-range*, specifically, the outcome selected maximizes the agent's utility less some cost that is a function of the outcome (e.g., this cost function can depend on other agents' reported preferences.).

The black-box reductions from Bayesian mechanism design to Bayesian algorithm design in the literature are based on obtaining an understanding of the distribution of outcomes produced by the algorithm through simulating the algorithm on samples from agents' preferences. Notice that, even for structurally simple problems, calculating the exact probability that a given outcome is selected by an algorithm can be #P-hard. For example, Hartline et al. (2015) show such a result for calculating the probability that a matching in a bipartite graph is optimal, for a simple explicitly given distribution of edge weights. On the other hand, a black-box reduction for mechanism design must produce exactly maximal-in-range outcomes merely from samples. This challenge motivates new questions for algorithm design from samples.

The Expectations from Samples Model. In traditional algorithm design, the inputs are specified to the algorithm exactly. In this paper, we formulate the *expectations from samples* model. This model calls for drawing an outcome from a distribution that is a precise function of the expectations of some random sources that are given only by sample access. Formally, a problem for this model is described by a function $f : [0,1]^n \to \Delta(X)$ where X is an abstract set of feasible outcomes and $\Delta(X)$ is the family of probability distributions over X. For any n input distributions on support $[0,1]$ with unknown expectations $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n)$, an algorithm for such a problem, with only sample access to each of the n input distributions, must produce sample outcome from X that is distributed exactly according to $f(\mu_1, \ldots, \mu_n)$.

Producing an outcome that is approximately drawn according to the desired distribution can typically be done from estimates of the expectations formed from sample averages (a.k.a., Monte Carlo sampling). On the other hand, exact implementation of many natural functions f is either impossible for information theoretic reasons or requires sophisticated techniques. Impossibility generally follows, for example, when f is discontinuous. The literature on *Bernoulli Factories* (e.g., Keane and O'Brien, 1994), which inspires our generalization to the expectations from samples model and provides some of the basic building blocks for our results, considers the special case where the input distribution and output distribution are both Bernoullis (i.e., supported on $\{0,1\}$).

We propose and solve two fundamental problems for the expectations from samples model. The

[1] A Bayesian algorithm is one that performs well in expectation when the input is drawn from a known distribution. By polynomial time, we mean polynomial in the number of agents and the combined "size" of their type spaces.

first problem considers the biases $\mathbf{p} = (p_1, \ldots, p_m)$ of m Bernoulli random variables as the marginal probabilities of a distribution on $\{1, \ldots, m\}$ (i.e., $\mathbf{p}$ satisfies $\sum_i p_i = 1$) and asks to sample from this distribution. We develop an algorithm that we call the Bernoulli Race to solve this problem.

The second problem corresponds to the "soft maximum" problem given by a regularizer that is a multiple $1/\lambda$ of the Shannon entropy function $H(\mathbf{p}) = -\sum_i p_i \log p_i$. The marginal probabilities on outcomes that maximize the expected value of the distribution over outcomes plus the entropy regularizer are given by exponential weights,[2] i.e., the function outputs i with probability proportional to $e^{\lambda p_i}$. A straightforward exponentiation and then reduction to the Bernoulli Race above does not have polynomial sample complexity. We develop an algorithm that we call the Fast Exponential Bernoulli Race to solve this problem.

Black-box Reductions in Mechanism Design. A special case of the problem that we must solve to apply the standard approach to black-box reductions is the *single-agent multiple-urns problem*. In this setting, a single agent faces a set X of urns, and each urn contains a random object whose distribution is unknown, but can be sampled. The agent's type determines his utility for each object; fixing this type, urn i is associated with a random real-valued reward with unknown expectation μ_i. Our goal is to allocate the agent his favorite urn, or close to it.

As described above, incentive compatibility requires an algorithm for selecting a high-value urn that is maximal-in-range. If we could exactly calculate the expected values $\mu_1, \ldots, \mu_n$ from the agent's type, this problem is trivial both algorithmically and from a mechanism design perspective: simply solicit the agent's type t then allocate him the urn with the maximum $\mu_i = \mu_i(t)$. As described above, with only sample access to the expected values of each urn, we cannot implement the exact maximum. Our solution is to apply the Fast Exponential Bernoulli Race as a solution to the regularized maximization problem in the expectations from samples model. This algorithm – with only sample access to the agent's values for each urn – will assign the agent to a random urn with a high expected value and is maximal-in-range.

The multi-agent reduction from Bayesian mechanism design to Bayesian algorithm design of Hartline et al. (2011, 2015) is based on solving a matching problem between multiple agents and outcomes, where an agent's value for an outcome is the expectation of a random variable which can be accessed only through sampling.[3] Specifically, this problem generalizes the above-described single-agent multiple-urns problem to the problem of matching agents to urns with the goal of approximately maximizing the total weight of the matching (the social welfare). Again, for incentive compatibility we require this expectations from samples algorithm to be maximal-in-range from each agent's perspective. Using methods from Agrawal and Devanur's (2015) work on stochastic online convex optimization, we reduce this matching problem to the single-agent multiple-urns problem.

As stated in the opening paragraph, our main result – obtained through the approach outlined above – is a polynomial time reduction from Bayesian incentive compatible mechanism design to Bayesian algorithm design. The analysis assumes that agents' values are normalized to the $[0, 1]$ interval and gives additive loss in the welfare. The reduction is an approximation scheme and the dependence of the runtime on the additive loss is inverse polynomial. The reduction depends polynomially on a suitable notion of the size of the space of agent preferences. For example, applied to environments where agents have preferences that lie in high-dimensional spaces, the runtime of the reduction depends polynomially on the number of points necessary to approximately cover each agent's space of preferences. More generally, the bounds we obtain are polynomial in the bounds of

[2]This is a standard relationship that has, for example, been employed in previous work in mechanism design (e.g., Huang and Kannan, 2012).

[3]Bei and Huang (2011) independently discovered a similar reduction based on solving a fractional assignment problem. Their reduction applies to finite, discrete type spaces and is approximately Bayesian incentive compatible.

Hartline et al. (2011, 2015) but the resulting mechanism, unlike in the proceeding work, is exactly Bayesian incentive compatible.

Organization. The organization of the paper separates the development of the expectations from samples model and its application to black-box reductions in Bayesian mechanism design. Section 2 introduces Bernoulli factories and reviews basic results from the literature. Section 3 defines two central problems in the expectations from samples model, sampling from outcomes with linear weights and sampling from outcomes with exponential weights, and gives algorithms for solving them. We return to mechanism design problems in Section 4 and solve the single-agent multiple urns problem. In Section 5 we give our main result, the reduction from Bayesian mechanism design to Bayesian algorithm design.

2 Basics of Bernoulli Factories

We use the terms *Bernoulli* and *coin* to refer to distributions over $\{1, 0\}$ and $\{\texttt{heads}, \texttt{tails}\}$, interchangeably. The Bernoulli factory problem is about generating new coins from old ones.

Definition 2.1 (Keane and O'Brien, 1994)**.** Given function $f : (0, 1) \to (0, 1)$, the *Bernoulli factory* problem is to output a sample of a Bernoulli variable with bias $f(p)$ (i.e. an $f(p)$-coin), given black-box access to independent samples of a Bernoulli distribution with bias $p \in (0, 1)$ (i.e. a p-coin).

To illustrate the Bernoulli factory model, consider the examples of $f(p) = p^2$ and $f(p) = e^{p-1}$. For the former one, it is enough to flip the p-coin twice and output 1 if both flips are 1, and 0 otherwise. For the latter one, the Bernoulli factory is still simple but more interesting: draw K from the Poisson distribution with parameter $\lambda = 1$, flip the p-coin K times and output 1 if all coin flips where 1, and 0 otherwise (see below).[4]

The question of characterizing functions f for which there is an algorithm for sampling $f(p)$-coins from p-coins has been the main subject of interest in this literature (Keane and O'Brien, 1994; Nacu and Peres, 2005). In particular, Keane and O'Brien (1994) provides necessary and sufficient conditions for f under which a Bernoulli factory exists. Moreover, Nacu and Peres (2005) suggests an algorithm for simulating an $f(p)$-coin based on polynomial envelopes of f. The canonical challenging problem of Bernoulli factories – and a primitive in the construction of more general Bernoulli factories – is the *Bernoulli Doubling* problem: $f(p) = 2p$ for $p \in (0, 1/2)$. See Łatuszyński (2010) for a survey on this topic.

Questions in Bernoulli factories can be generalized to multiple input coins. Given $f : (0, 1)^m \to (0, 1)$, the goal is sample from a Bernoulli with bias $f(p_1, \ldots, p_m)$ given sample access to m independent Bernoulli variables with unknown biases $\mathbf{p} = (p_1, \ldots, p_m)$. Linear functions f were studied and solved by Huber (2015). For example, the special case $m = 2$ and $f(p_1, p_2) = p_1 + p_2$, a.k.a., *Bernoulli Addition*, can be solved by reduction to the Bernoulli Doubling problem (formalized below).

Questions in Bernoulli factories can be generalized to allow input distributions over real numbers on the unit interval $[0, 1]$ (rather than Bernoullis over $\{0, 1\}$). In this generalization the question is to produce a Bernoulli with bias $f(\mu)$ with sample access to draws from a distribution supported on $[0, 1]$ with expectation μ. These problems can be easily solved by reduction to the Bernoulli factory problem:

0. *Continuous to Bernoulli*: Can implement Bernoulli with bias μ with one sample from distribution $\mathcal{D}$ with expectation μ. Algorithm:

[4]The Poisson distribution with parameter λ has probability of $K = k$ as $\lambda^k e^{-\lambda}/k!$.

- Draw $Z \sim \mathcal{D}$ and $P \sim \text{Bern}[Z]$.
- Output P.

Below are enumerated the important building blocks for Bernoulli factories.

1. *Bernoulli Down Scaling*: Can implement $f(p) = \lambda \cdot p$ for $\lambda \in [0, 1]$ with one sample from $\text{Bern}[p]$. Algorithm:
   - Draw $\Lambda \sim \text{Bern}[\lambda]$ and $P \sim \text{Bern}[p]$.
   - Output $\Lambda \cdot P$ (i.e., 1 if both coins are 1, otherwise 0).
2. *Bernoulli Doubling:* Can implement $f(p) = 2p$ for $p \in (0, 1/2 - \delta]$ with $O(1/\delta)$ samples from $\text{Bern}[p]$ in expectation. The algorithm is complicated, see Nacu and Peres (2005).
3. *Bernoulli Probability Generating Function:* Can implement $f(p) = \mathbf{E}_{K \sim \mathcal{D}}\left[p^K\right]$ for distribution $\mathcal{D}$ over non-negative integers with $\mathbf{E}_{K \sim \mathcal{D}}[K]$ samples from $\text{Bern}[p]$ in expectation. Algorithm:
   - Draw $K \sim \mathcal{D}$ and $P_1, \ldots, P_K \sim \text{Bern}[p]$ (i.e., K samples).
   - Output $\prod_i P_i$ (i.e., 1 if all K coins are 1, otherwise 0).
4. *Bernoulli Exponentiation:* Can implement $f(p) = \exp(\lambda(p-1))$ for $p \in [0, 1]$ and non-negative constant λ with λ samples from $\text{Bern}[p]$ in expectation. Algorithm: Apply the Bernoulli Probability Generating Function algorithm for the Poisson distribution with parameter λ.
5. *Bernoulli Averaging:* Can implement $f(p_1, p_2) = (p_1 + p_2)/2$ with one sample from $\text{Bern}[p_1]$ or $\text{Bern}[p_2]$. Algorithm:
   - Draw $Z \sim \text{Bern}[1/2]$, $P_1 \sim \text{Bern}[p_1]$, and $P_2 \sim \text{Bern}[p_2]$.
   - Output P_{Z+1}.
6. *Bernoulli Addition:* Can implement $f(p_1, p_2) = p_1 + p_2$ for $p_1 + p_2 \in [0, 1 - \delta]$ with $O(1/\delta)$ samples from $\text{Bern}[p_1]$ and $\text{Bern}[p_2]$ in expectation. Algorithm: Apply Bernoulli Doubling to Bernoulli Averaging.

It may seem counterintuitive that Bernoulli Doubling is much more challenging that Bernoulli Down Scaling. Notice, however, that for a coin with bias $p = 1/2$, Bernoulli Doubling with a finite number of coin flips is impossible. The doubled coin must be deterministically heads, while any finite sequence of coin flips of $\text{Bern}[1/2]$ has non-zero probability of occuring. On the other hand a coin with probability $p = 1/2 - \delta$ for some small δ has a similar probability of each sequence but Bernoulli Doubling must sometimes output tails. Thus, Bernoulli Doubling must require a number of coin flips that goes to infinity as δ goes to zero.

3 The Expectations from Samples Model

The expectations from samples model is a combinatorial generalization of the Bernoulli factory problem. The goal is to select an outcome from a distribution that is a function of the expectations of a set of input distributions. These input distributions can be accessed only by sampling.

Definition 3.1. Given function $f : (0, 1)^n \to \Delta(X)$ for domain X, the *expectations from samples* problem is to output a sample from $f(\boldsymbol{\mu})$ given black-box access to independent samples from n distributions supported on $[0, 1]$ with expectations $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n) \in (0, 1)^n$.

Without loss of generality, by the Continuous to Bernoulli construction of Section 2, the input random variables can be assumed to be Bernoullis and, thus, this expectations of samples model can be viewed as a generalization of the Bernoulli factory question to output spaces X beyond $\{0, 1\}$. In this section we propose and solve two fundamental problems for the expectations of samples model. In these problems the outcomes are the a finite set of m outcomes $X = \{1, \ldots, m\}$ and the input distributions are m Bernoulli distributions with biases $\mathbf{p} = (p_1, \ldots, p_m)$.

In the first problem, biases correspond to the marginal probabilities with which each of the outcomes should be selected. The goal is to produce random i from X so that the probability of i is exactly its marginal probability p_i. More generally, if the biases do not sum to one, this problem is equivalently the problem of *random selection with linear weights*.

The second problem we solve corresponds to a regularized maximization problem, or specifically *random selection from exponential weights*. For this problem the baiases of the m Bernoulli input distributions correspond to the weights of the outcomes. The goal is to produce a random i from X according to the distribution given by exponential weights, i.e., the probability of selecting i from X is $e^{\lambda p_i} / \sum_j e^{\lambda p_j}$.

3.1 Random Selection with Linear Weights

Definition 3.2 (**Random Selection with Linear Weights**)**.** The *random selection with linear weights* problem is to sample from the probability distribution $f(\mathbf{v})$ defined by $\mathbf{Pr}_{I \sim f(\mathbf{v})}[I = i] = v_i / \sum_j v_j$ for each i in $\{1, \ldots, m\}$ with only sample access to distributions with expectations $\mathbf{v} = (v_1, \ldots, v_m)$.

We solve the random selection with linear weights problem by an algorithm that we call the *Bernoulli race* (Algorithm 1). The algorithm repeatedly picks a coin uniformly at random and flips it. The winning coin is the first one to come up heads in this process.

Algorithm 1 Bernoulli Race

1: **input** sample access to m coins with biases $v_1, \ldots, v_m$.
2: **loop**
3: Draw I uniformly from $\{1, \ldots, m\}$ and draw P from input distribution I.
4: If P is `heads` then output I and halt.
5: **end loop**

Theorem 3.1. *The Bernoulli Race (Algorithm 1) samples with linear weights (Definition 3.2) with an expected $m / \sum_i v_i$ samples from input distributions with biases $v_1, \ldots, v_n$.*

Proof. At each iteration, the algorithm terminates if the flipped coin outputs 1 and iterates otherwise. Since the coin is chosen uniformly at random, the probability of termination at each iteration is $\frac{1}{m} \sum_i v_i$. The total number of iterations (and number of samples) is therefore a geometric random variable with expectation $m / \sum_i v_i$.

The selected outcome also follows the desired distribution, as shown below.

$$\begin{aligned}
\mathbf{Pr}[i \text{ is selected}] &= \sum_{k=1}^{\infty} \mathbf{Pr}[i \text{ is selected at time } k] \; \mathbf{Pr}[\text{algorithm reaches time } k] \\
&= \frac{v_i}{m} \sum_{k=1}^{\infty} \left(1 - \frac{1}{m} \sum\nolimits_j v_j\right)^{k-1} = \frac{\frac{v_i}{m}}{\frac{1}{m} \sum_j v_j} = \frac{v_i}{\sum_j v_j}.
\end{aligned}$$

□

3.2 Random Selection with Exponential Weights

Definition 3.3 (**Random Selection with Exponential Weights**). For parameter $\lambda > 0$, the *random selection with exponential weights* problem is to sample from the probability distribution $f(\mathbf{v})$ defined by $\mathbf{Pr}_{I \sim f(\mathbf{v})}[I = i] = \exp(\lambda v_i) / \sum_j \exp(\lambda v_j)$ for each i in $\{1, \ldots, m\}$ with only sample access to distributions with expectations $\mathbf{v} = (v_1, \ldots, v_m)$.

The *Basic Exponential Bernoulli Race*, below, samples from the exponential weights distribution. The algorithm follows the paradigm of picking one of the input distributions, exponentiating it, sampling from the exponentiated distribution, and repeating until one comes up heads. While this algorithm does not generally run in polynomial time, it is a building block for one that does.

Algorithm 2 The Basic Exponential Bernoulli Race (with parameter $\lambda > 0$)

1: **input** Sample access to m coins with biases $v_1, \ldots, v_m$.
2: For each i, apply Bernoulli Exponentiation to coin i to produce coin with bias $\tilde{v}_i = \exp(\lambda(v_i - 1))$.
3: Run the Bernoulli Race on the coins with biases $\tilde{\mathbf{v}} = (\tilde{v}_1, \ldots, \tilde{v}_m)$.

Theorem 3.2. *The Basic Exponential Bernoulli Race (Algorithm 2) samples with exponential weights (Definition 3.3) with an expected $\lambda m e^{\lambda(1-v_{\max})}$ samples from input distributions with biases $v_1, \ldots, v_n$ and $v_{\max} = \max_i v_i$.*

Proof. The correctness and runtime follows from the correctness and runtimes of Bernoulli Exponentiation and the Bernoulli Race. □

3.3 The Fast Exponential Bernoulli Race

Sampling from exponential weights is typically used as a "soft maximum" where the parameter λ controls how close the selected outcome is to the true maximum. For such an application, exponential dependence on λ in the runtime would be prohibitive. Unfortunately, when $v_{\max}$ is bounded away from one, the runtime of the Basic Logistic Bernoulli Race (Algorithm 2; Theorem 3.2) is exponential in λ. A simple observation allows allows the resolution of this issue: the exponential weights distribution is invariant to any uniform additive shift of all weights. This section applies this idea to develop the *Fast Logistic Bernoulli Race*.

Observe that for any given parameter ϵ, we can easily implement a Bernoulli random variable Z whose bias z is within an additive ϵ of $v_{\max}$. Note that, unlike the other algorithms in this section, a precise relationship between z and $v_1, \ldots, v_m$ is not required.

Lemma 3.3. *For parameter $\epsilon \in (0, 1]$, there is an algorithm for sampling from a Bernoulli random variable with bias $z \in [v_{\max} - \epsilon, v_{\max} + \epsilon]$, where $v_{\max} = \max_i v_i$, with $O(\frac{m}{\epsilon^2} \cdot \log(\frac{m}{\epsilon}))$ samples from input distributions with biases $v_1, \ldots, v_m$.*

Proof. The algorithm is as follows: Sample $\frac{4}{\epsilon^2} \log(\frac{4m}{\epsilon})$ times from each of the m coins, let $\hat{v}_i$ be the empirical estimate of coin i's bias obtained by averaging, then apply the Continuous to Bernoulli algorithm (Section 2) to map $\hat{v}_{\max} = \max_i \hat{v}_i$ to a Bernoulli random variable.

Standard tail bounds imply that $|\hat{v}_{\max} - v_{\max}| < \epsilon/2$ with probability at least $1 - \epsilon/2$, and therefore $z = \mathbf{E}[\hat{v}_{\max}] \in [v_{\max} - \epsilon, v_{\max} + \epsilon]$. □

Since we are interested in a fast logistic Bernoulli race as λ grows large, we restrict attention to $\lambda > 4$. We set $\epsilon = 1/\lambda$ in the estimation of $v_{\max}$ (by Lemma 3.3). This estimate will be used to boost the bias of each distribution in the input so that the maximum bias is at least $1 - 3\epsilon$.

The boosting of the bias is implemented with Bernoulli Addition which, to be fast, requires the cumulative bias be bounded away from one. Thus, the probabilities are scaled down by a factor of $1 - 2\epsilon$, this scaling is subsequently counterbalanced by adjusting the parameter λ. The formal details are given below.

Algorithm 3 Fast Exponential Bernoulli Race (with parameter $\lambda > 4$)

1: **input** Sample access to m coins with biases $v_1, \ldots, v_m$.
2: Let $\epsilon = 1/\lambda$.
3: Construct a coin with bias $z \in [v_{\max} - \epsilon, v_{\max} + \epsilon]$ (from Lemma 3.3).
4: Apply Bernoulli Down Scaling to a coin with bias $1 - z$ to implement a coin with bias $(1 - 2\epsilon)(1 - z)$.
5: For all i, apply Bernoulli Down Scaling to implement a coin with bias $(1 - 2\epsilon)v_i$.
6: For all i, apply Bernoulli Addition to implement coin with bias $v'_i = (1 - 2\epsilon)v_i + (1 - 2\epsilon)(1 - z)$.
7: Run the Basic Exponential Bernoulli Race with parameter $\lambda' = \frac{\lambda}{1-2\epsilon}$ on the coins with bias $v'_1, \ldots, v'_m$.

Theorem 3.4. *The Fast Exponential Bernoulli Race (Algorithm 3) samples with exponential weights (Definition 3.3) with an expected $O(\lambda^4 m^2 \log(\lambda m))$ samples from the input distributions.*

Proof. The correctness and runtime follows from the correctness and runtimes of the Basic Exponential Bernoulli Race, Bernoulli Doubling, Lemma 3.3 (for estimate of $v_{\max}$), and the facts that $\lambda' v'_i = \lambda(v_i + 1 - z)$ and that the distribution given by exponential weights is invariant to additive shifts of all weights.

A detailed analysis of the runtime follows. Since the algorithm builds a number of sampling subroutines in a hierarchy, we analyze the runtime of the algorithm and the various subroutines in a bottom up fashion. Steps 3 and 4 implement a coin with bias $(1 - 2\epsilon)(1 - z)$ with runtime $O(\lambda^2 m \cdot \log(\lambda m))$ per sample, as per the bound of Lemma 3.3. The coin implemented in Step 5 is sampled in constant time. Observe that $v'_i \leq (1 - 2\epsilon)(1 + v_i - v_{\max} + \epsilon) \leq 1 - \epsilon$, and the runtime of Bernoulli Doubling implies that $O(\lambda)$ samples from the coins of Steps 4 and 5 suffice for sampling $\mathrm{Bern}[v'_i]$; we conclude that a v'_i-coin can be sampled in time $O(\lambda^3 m \cdot \log(\lambda m))$. Finally, note that for $v'_{\max} = \max_i v'_i$, we have $v'_{\max} \geq 1 - 3\epsilon$; Theorem 3.2 then implies that the Basic Exponential Bernoulli Race samples at most $\lambda' m e^{\lambda' 3\epsilon} \leq 2e^6 \lambda m = O(\lambda m)$ times from the $\mathbf{v}'$-coins; we conclude the claimed runtime. □

4 The Single-Agent Multiple-Urns Problem

We investigate incentive compatible mechanism design for the *single-agent multiple-urns* problem. Informally, mechanism is needed to assign an agent to one of many urns. Each urn contains objects and the agent's value for being assigned to an urn is taken in expectation over objects from the urn. The problem asks for an incentive compatible mechanism with good welfare (i.e., the value of the agent for the assigned urn).

4.1 Problem Definition and Notations

A single agent with type t from type space $\mathcal{T}$ desires an object o from outcome space $\mathcal{O}$. The agent's value for an outcome o is a function of her type t and denoted by $v(t, o) \in [0, 1]$. The agent is a risk-neutral quasi-linear utility maximizer with utility $\mathbf{E}_o[v(t, o)] - p$ for randomized outcome o and expected payment p. There are m urns. Each urn j is given by a distribution $\mathcal{D}_j$ over outcomes in $\mathcal{O}$. If the agent is assigned to urn j she obtains an object from the urn's distribution $\mathcal{D}_j$.

A mechanism can solicit the type of the agent (who may misreport if she desires). We further assume (1) the mechanism has black-box access to evaluate $v(t, o)$ for any type t and outcome o, (2) the mechanism has sample access to the distribution $\mathcal{D}_j$ of each urn j. The mechanism may draw objects from urns and evaluate the agent's reported value for these objects, but then must ultimately assign the agent to a single urn and charge the agent a payment. The urn and payment that the agent is assigned are random variables in the mechanism's internal randomization and randomness from the mechanisms potential samples from the urns' distributions.

The distribution of the urn the mechanism assigns to an agent, as a function of her type t, is denoted by $\mathbf{x}(t) = (x_1(t), \ldots, x_m(t))$ where $x_j(t)$ is the marginal probability that the agent is assigned to urn j. Denote the expected value of the agent for urn j by $v_j(t) = \mathbf{E}_{o \sim \mathcal{D}_j}[v(t, o)]$. The expected welfare of the mechanism is $\sum_j v_j(t)\, x_j(t)$. The expected payment of this agent is denoted by $p(t)$. The agent's utility for the outcome and payment of the mechanism is given by $\sum_j v_j(t)\, x_j(t) - p(t)$. Incentive compatibility is defined by the agent with type t preferring her outcome and payment to that assigned to another type t'.

Definition 4.1. A single-agent mechanism $(\mathbf{x}, p)$ is *incentive compatible* if, for all $t, t' \in \mathcal{T}$:

$$\sum\nolimits_j v_j(t)\, x_j(t) - p(t) \geq \sum\nolimits_j v_j(t)\, x_j(t') - p(t') \tag{1}$$

A multi-agent mechanism is Bayesian Incentive Compatible (BIC) if equation (1) holds for the outcome of the mechanism in expectation of the truthful reports of the other agents.

4.2 Incentive Compatible Approximate Scheme

If the agent's expected value for each urn is known, or equivalently mechanism designer knows the distributions $\mathcal{D}_j$ for all urns j rather than only sample access, this problem would be easy and admits a trivial optimal mechanism: simply select the urn maximizing the agent's expected value $v_j(t)$ according to her reported type t, and charge her a payment of zero. What makes this problem interesting is that the designer is restricted to only *sample* the agent's value for an urn. In this case, the following Monte-carlo adaptation of the trivial mechanism is tempting: sample from each urn sufficiently many times to obtain a close estimate $\tilde{v}_j(t)$ of $v_j(t)$ with high probability (up to any desired precision $\delta > 0$), then choose the urn j maximizing $\tilde{v}_j(t)$ and charge a payment of zero. This mechanism is not incentive compatible, as illustrated by a simple example.

Example Consider two urns. Urn A contains only outcome o_2, whereas B two contains a mixture of outcomes o_1 and o_3, with o_1 slightly more likely than o_3. Now consider an agent who has (true) values 1, 2, and 3 for outcomes o_1, o_2, and o_3 respectively. If this agent reports her true type, the trivial Monte-carlo mechanism — instantiated with any desired finite degree of precision — assigns her urn A most of the time, but assigns her urn B with some nonzero probability. The agent gains by misreporting her value of outcome o_3 as 0, since this guarantees her preferred urn A.

The above example might seem counter-intuitive, since the trivial Monte-carlo mechanism appears to be doing its best to maximize the agent's utility, up to the limits of (unavoidable) sampling error. One intuitive rationalization is the following: an agent can slightly gain by procuring (by whatever means) more precise information about the distributions $\mathcal{D}_j$ than that available to the mechanism, and using this information to guide her strategic misreporting of her type. This raises the following question:

Question: *Is there an incentive-compatible mechanism for the single-agent multiple-urns problem which achieves welfare within ϵ of the optimal, and samples only poly$(m, \frac{1}{\epsilon})$ times (in expectation) from the urns?*

We resolve the above question in the affirmative. We present approximation scheme for this problem that is based on our solution to the problem of random selection with exponential weights (Section 3.2). The solution to the single-agent multiple-urns problem is a main ingredient in the Bayesian mechanism that we propose in Section 5 as our black-box reduction mechanism.

To explain the approximate scheme, we start by recalling the following standard theorem in mechanism design.

Theorem 4.1. *For outcome rule $\mathbf{x}$, there exists payment rule p so that single-agent mechanism $(\mathbf{x}, p)$ is incentive compatible if and only if $\mathbf{x}$ is maximal in range, i.e., $\mathbf{x}(t) \in \operatorname{argmax}_{\mathbf{x}'} \sum_j v_j(t)\, x'_j - c(\mathbf{x}')$, for some cost function $c(\cdot)$.*[5]

The payments that satisfy Theorem 4.1 can be easily calculated with black-box access to outcome rule $\mathbf{x}(\cdot)$. For a single-agent problem, this payment can be calculated in two calls to the function $\mathbf{x}(\cdot)$, one on the agent's reported type t and the other on a type randomly drawn from the path between the origin and t. Further discussion and details are given in Appendix A. It suffices, therefore, to identify a mechanism that samples from urns and assigns the agent to an urn that induces an outcome rule $\mathbf{x}(\cdot)$ that is good for welfare, i.e., $\sum_i v_j(t)\, x_j(t)$, and is maximal in range. The following theorem solves the problem.

Theorem 4.2. *There is an incentive-compatible mechanism for the single-agent multiple-urns problem which achieves an additive ϵ-approximation to the optimal welfare in expectation, and runs in time $O(m^2(\frac{\log m}{\epsilon})^5)$ in expectation.*

Proof. Consider the problem of selecting a distribution over urns to optimize welfare plus (a scaling of) the Shannon entropy function, i.e., $\mathbf{x}(t) = \operatorname{argmax}_{\mathbf{x}'} v_j(t)\, x'_j - (1/\lambda) \sum_j x'_j \log x'_j$.[6] It is well known that the optimizer $\mathbf{x}(t)$ is given by exponential weights, i.e., the marginal probability of assigning the jth urn is given by $x_j(t) = \exp(\lambda v_j(t)) / \sum_{j'} \exp(\lambda v_{j'}(t))$. In Section 3.3 we gave a polynomal time algorithm for sampling from exponential weights, specifically, the Fast Exponential Bernoulli Race (Algorithm 3). Proper choice of the parameter λ trades off faster runimes with increased welfare loss due to entropy term. The entropy is maximized at the uniform distribution $\mathbf{x}' = (1/m, \ldots, 1/m)$ with entropy $\log m$. Thus, choosing $\lambda = \log m/\epsilon$ guarantees that the welfare is within an additive ϵ of the optimal welfare $\max_j v_j(t)$. The bound of the theorem then follows from the analysis of the Fast Exponential Bernoulli Race (Theorem 3.4) with this choice of λ. □

5 A Bayesian Incentive Compatible Black-box Reduction

A central question at the interface between algorithms and economics is on the existence of black-box reductions for mechanism design. Given black-box access to any algorithm that maps inputs to outcomes, can a mechanism be constructed that (a) induces agents to truthfully report the inputs and (b) produces an outcome that is as good as the one produced by the algorithm? The mechanism must be computationally tractable, specifically, making no more than a polynomial number of elementary operations and black-box calls to the algorithm.

[5] The "only if" direction of this theorem requires that the type space $\mathcal{T}$ be rich enough so that the induced space of values across the urns is $\{(v_1(t), \ldots, v_m(t)) : t \in \mathcal{T}\} = [0, 1]^m$.

[6] The additive entropy term can be interpreted as a negative cost vis-à-vis Theorem 4.1.

A line of research initiated by Hartline and Lucier (2010, 2015) demonstrated that, for the welfare objective, Bayesian black-box reductions exist. In the Bayesian setting, agents' types are drawn from a distribution. The algorithm is assumed to obtain good welfare for types from this distribution. The constructed mechanism is an approximation scheme; For any ϵ it gives a mechanism that is Bayesian incentive compatible (Definition 4.1) and obtains a welfare that is at least an additive ϵ form the algorithms welfare. Before formalizing this problem, for further details on Bayesian mechanism design and our set of notations in this paper, which are based on those in Hartline et al. (2015), we refer the reader to Appendix B.

Definition 5.1 (BIC black-box reduction problem)**.** Given black-box oracle access to an allocation algorithm $\mathcal{A}$, construct an allocation algorithm $\tilde{\mathcal{A}}$ that is Bayesian incentive compatible; approximately preserves welfare, i.e., any agent's expected welfare under $\tilde{\mathcal{A}}$ is at least that under $\mathcal{A}$ less ϵ; and runs in polynomial time in n and $1/\epsilon$.

In this literature, Hartline and Lucier (2010, 2015) solve the case of single-dimensional agents and Hartline et al. (2011, 2015) solve the case of multi-dimensional agents with discrete type spaces. For the relaxation of the problem where only approximate incentive compatibility is required, Bei and Huang (2011) solve the case of multi-dimensional agents with discrete type space, and Hartline et al. (2011, 2015) solve the general case. These reductions are approximation schemes that are polynomial in the number of agents, the desired approximation factor, and a measure of the size of the agents' type spaces (e.g., its dimension).

5.1 Surrogate Selection and the Replica-Surrogate Matching

A main conclusion of the literature on Bayesian reductions for mechanism design is that the multi-agent problem of reducing Bayesian mechanism design to algorithm design, itself, reduces to a single-agent problem of *surrogate selection*. Consider any agent in the original problem and the *induced algorithm* with the inputs form other agents hardcoded as random draws from their respective type distributions. The induced algorithm maps the type of this agent to a distribution over outcomes. If this distribution over outcomes is maximal-in-range then there exists payments for which the induced algorithm is incentive compatible (Theorem 4.1). If not, the problem of surrogate selection is to map the type of the agent to an input to the algorithm to satisfy three properties:

(a) The composition of surrogate selection and the induced algorithm is maximal-in-range,

(b) The composition approximately preserves welfare,

(c) The surrogate selection preserves the type distribution.

Condition (c), a.k.a. *stationarity*, implies that fixing the non-maximaility-of-range of the algorithm for a particular agent does not affect the outcome for any other agents. With such an approach each agent's incentive problem can be resolved independently from that of other agents.

Theorem 5.1 (Hartline et al., 2015)**.** *The composition of an algorithm with a profile of surrogate selection rules satisfying conditions (a)–(c) is Bayesian incentive compatible and approximately preserves the algorithm's welfare (the loss in welfare is the sum of the losses in welfare of each surrogate selection rule).*

The surrogate selection rule of Hartline et al. (2015) is based on setting up a matching problem between random types from the distribution (replicas) and the outcomes of the algorithm on random types from the distribution (surrogates). The true type of the agent is one of the replicas, and

the surrogate selection rule outputs the surrogate to which this replica is matched. This approach addresses the three properties of surrogate selection rules as (a) if the matching selected is maximal-in-range then the composition of the surrogate selection rule with the induced algorithm is maximal-in-range, (b) the welfare of the matching is the welfare of the reduction and the optimal matching approximates the welfare of the original algorithm, and (c) any maximal matching gives a stationary surrogate selection rule. For a detailed discussion on why maximal-in-range matching will result in a BIC mechanism after composing the corresponding surrogate selection rule with the allocation algorithm, we refer the interested reader to Lemma C.1 and Lemma C.2 in Appendix C.

Definition 5.2. The *replica-surrogate matching* surrogate selection rule; for a k-to-1 matching algorithm M, a integer market size m, and load k; maps a type t to a surrogate type as follows:

1. Pick the real-agent index i^* uniformly at random from $\{1, \ldots, km\}$.
2. Define the *replica type profile* $\mathbf{r}$, an km-tuple of types by setting $r_{i^*} = t$ and sampling the remaining $km - 1$ replica types $\mathbf{r}_{-i^*}$ i.i.d. from the type distribution F.
3. Sample the *surrogate type profile* $\mathbf{s}$, an m-tuple of i.i.d. samples from the type distribution F.
4. Run matching algorithm M on the complete bipartite graph between replicas and surrogates.
5. Output the surrogate j^* that is matched to i^*.

The value that a replica obtains for the outcome that the induced algorithm produces for a surrogate, henceforth, *surrogate outcome*, is a random variable. The analysis of Hartline et al. (2015) is based on the study of an ideal computational model where the value of any replica for any surrogate outcome $\mathbf{E}_{o \sim \mathcal{A}(s_j)}[v(r_i, o)]$ is known exactly. In this computationally-unrealistic model and with these values as weights, the maximum weight matching algorithm can be employed in the replica-surrogate matching surrogate selection rule above, and it results in a Bayesian incentive compatible mechanism. Hartline et al. (2015) analyze the welfare of the resulting mechanism in the case where the load is $k = 1$, prove that conditions (a)-(c) are satisfied, and give (polynomial) bounds on the size m that is necessary for the expected welfare of the mechanism to be an additive ϵ from that of the algorithm.

Remark Given a BIC allocation algorithm $\tilde{\mathcal{A}}$ through a replica-surrogate matching surrogate selection, the payments that satisfy Bayesian incentive compatibility can be easily calculated with black-box access to $\tilde{\mathcal{A}}$ (see Appendix A).

If M is maximum matching, conditions (a)-(c) clearly continue to hold for our generalization to load $k > 1$. Moreover, the welfare of the reduction is monotone non-decreasing in k.

Lemma 5.2. *In the ideal computational model (where the value of a replica for being matched to a surrogate is given exactly) the per-replica welfare of the replica-surrogate maximum matching is monotone non-decreasing in load k.*

Proof. Consider a non-optimal matching that groups replicas into k groups of size m and finds the optimal 1-to-1 matching between replicas in each group and the surrogates. As these are random $(k = 1)-$matchings, the expected welfare of each such matching is equal to the expected welfare of the $(k = 1)-$matching. These matchings combine to give a feasible matching between the mk replicas and m surrogates. The total expected welfare of the optimal k-to-1 matching between replicas and surrogates is no less than k times the expected welfare of the $(k = 1)-$matching. Thus, the per-replica welfare, i.e., normalized by mk, is monotone in k. □

Our main result is an approximation scheme for the ideal reduction of Hartline et al. (2015). We identify a $k > 1$ and a polynomial time (in m and $1/\epsilon$) k-to-1 matching algorithm for the black-box model and prove that the expected welfare of this matching algorithm (per-replica) is within an additive ϵ of the expected welfare per-replica of the optimal matching in the ideal model with load $k = 1$ (as analyzed by Hartline et al., 2015). The welfare of the ideal model is monotone non-decreasing in load k (Lemma 5.2); therefore it will be sufficient to identify a polynomial load k where there is a polynomial time algorithm in the black-box model that has ϵ loss relative to the ideal model for that same load k.

In the remainder of this section we replace this ideal matching algorithm with an approximation scheme for the black-box model where replica values for surrogate outcomes can only be estimated by sampling. For any ϵ our algorithm gives an ϵ additive loss of the welfare of the ideal algorithm with only a polynomial increase to the runtime. Moreover, the algorithm produces a perfect (and so maximal) matching, and therefore the surrogate selection rule is stationary; and the algorithm is maximal-in-range for any replica (including the true type of the agent), and therefore the resulting mechanism is Bayesian incentive compatible.

5.2 Entropy Regularized Matching

In this section we define an entropy regularized bipartite matching problem and discuss its solution. We will refer to the left-hand-side vertices as replicas and the right-hand-side vertices as surrogates. The weights on the edge between replica $i \in \{1, \ldots, km\}$ and surrogate $j \in \{1, \ldots, m\}$ will be denoted by $v_{i,j}$. In our application to the replica-surrogate matching defined in the previous section, the weights will be set to $v_{i,j} = \mathbf{E}_{o \sim \mathcal{A}(s_j)}[v(r_i, o)]$ for $(i, j) \in [km] \times [m]$.

Definition 5.3. For weights $\mathbf{v} = [v_{i,j}]_{(i,j) \in [km] \times [m]}$, the entropy regularized matching program for parameter $\delta > 0$ is:

$$\begin{aligned} \max_{\{x_{i,j}\}_{(i,j) \in [km] \times [m]}} & \sum\nolimits_{i,j} x_{i,j}\, v_{i,j} - \delta \sum\nolimits_{i,j} x_{i,j} \log x_{i,j}, \\ \text{s.t.} \quad & \sum\nolimits_{i} x_{i,j} \le k && \forall j \in [m], \\ & \sum\nolimits_{j} x_{i,j} \le 1 && \forall i \in [km]. \end{aligned}$$

The optimal value of this program is denoted $\mathrm{OPT}(\mathbf{v})$.

The dual variables for right-hand-side constraints of the matching polytope can be interpreted as *prices* for the surrogate outcomes. Given prices, the *utility* of a replica for a surrogate outcome given prices is the difference between the replica's value and the price. The following lemma shows that for the right choice of dual variables, the maximizer of the entropy regularized matching program is given by exponential weights with weights equal to the utilities.

Observation 1. *For the optimal Lagrangian dual variables $\boldsymbol{\alpha}^* \in \mathbb{R}^m$ for surrogate feasibility in the entropy regularized matching program (Definition 5.3), namely,*

$$\boldsymbol{\alpha}^* = \mathrm{argmin}_{\boldsymbol{\alpha} \ge \mathbf{0}} \max_{\mathbf{x}} \left\{ \mathcal{L}(\mathbf{x}, \boldsymbol{\alpha}) : \sum\nolimits_{j} x_{i,j} \le 1 \ , \ \forall i \right\}$$

where $\mathcal{L}(\mathbf{x}, \boldsymbol{\alpha}) \triangleq \sum_{i,j} x_{i,j}\, v_{i,j} - \delta \sum_{i,j} x_{i,j} \log x_{i,j} + \sum_j \alpha_j (k - \sum_i x_{i,j})$ is the Lagrangian objective function; the optimal solution $\mathbf{x}^*$ to the primal is given by exponential weights:

$$x^*_{i,j} = \frac{\exp\left(\frac{v_{i,j} - \alpha^*_j}{\delta}\right)}{\sum_{j'} \exp\left(\frac{v_{i,j'} - \alpha^*_{j'}}{\delta}\right)}, \quad \forall i, j.$$

Observation 1 recasts the entropy regularized matching as, for each replica, sampling from the distribution of exponential weights. For any replica i and fixed dual variables $\boldsymbol{\alpha}$ our Fast Exponential Bernoulli Race (Algorithm 3) gives a polynomial time algorithm for sampling from the distribution of exponential weights in the expectations from samples computational model.

Lemma 5.3. *For replica i and any prices (dual variables) $\boldsymbol{\alpha} \in [0,h]^m$, allocating a surrogate j drawn from the exponential weights distribution*

$$x_{i,j} = \frac{\exp\left(\frac{v_{i,j}-\alpha_j}{\delta}\right)}{\sum_{j'} \exp\left(\frac{v_{i,j'}-\alpha_{j'}}{\delta}\right)}, \quad \forall j \in [m], \tag{2}$$

is maximal-in-range, as defined in Definition 4.1, and this random surrogate j can be sampled with $O\left(\frac{h^4 m^2 \log(hm/\delta)}{\delta^4}\right)$ samples from replica-surrogate-outcome value distributions.

Proof. To see that the distribution is maximal-in-range when assigning surrogate outcome j to replica i, consider the regularized welfare maximization

$$\text{argmax}_{\mathbf{x}'} \sum_j v_{i,j}\, x'_j - \delta \sum_j x'_j \log x'_j - \sum_j \alpha_j x'_j$$

for replica i. Similar to Observation 1, first-order conditions imply that the exponential weights distribution in (2) is the unique maximizer of this concave program.

To apply the Fast Exponential Bernoulli Race to the utilities, of the form $v_{i,j} - \alpha_j \in [-h, 1]$, we must first normalize them to be on the interval $[0,1]$. This normalization is accomplished by adding h to the utilities (which has no effect on the exponential weights distribution, and therefore preserves maximality-in-range), and then scaling by $1/(h+1)$. The scaling needs to be corrected by setting λ in the Fast Exponential Bernoulli Race (Algorithm 3) to $(h+1)/\delta$. The expected number of samples from the value distributions that are required by the algorithm, per Theorem 3.4, is $O(h^4 m^2 \log(hm/\delta)\delta^{-4})$.

□

If we knew the optimal Lagrangian variables $\boldsymbol{\alpha}^*$ from Observation 1, it would be sufficient to define the surrogate selection rule by simply sampling from the exponential weights distribution (which is polynomial time per Lemma 5.3) that corresponds to the agent's true type (indexed i^*). Notice that the wrong values of $\boldsymbol{\alpha}$ correspond to violating primal constraints (for the surrogates) and thus the outcome from sampling from exponential weights for such α would not correspond to a maximal-in-range matching. In the next section we give a polynomial time approximation scheme that is maximal-in-range for each replica and approximates sampling from the correct $\boldsymbol{\alpha}^*$.

5.3 Online Entropy Regularized Matching

In this section, we reduce the entropy regularized matching problem to the problem of sampling from exponential weights (as described in Lemma 5.3) via an online algorithm. Consider replicas being drawn adversarially, but in a random order, over times $1, \ldots, km$. The basic observation is that approximate dual variables $\boldsymbol{\alpha}$ are sufficient for an online assignment of each replica to a surrogate via Lemma 5.3 to approximate the optimal (offline) regularized matching. Recall, the replicas are independently and identically distributed in the original problem.

Our construction borrows techniques used in designing online algorithms for stochastic online convex programming problems (Agrawal and Devanur, 2015; Chen and Wang, 2013), and stochastic online packing problems (Agrawal et al., 2009; Devanur et al., 2011; Badanidiyuru et al., 2013;

Kesselheim et al., 2014). Our online algorithm (Algorithm 4, below) considers the replicas in order, updates the dual variables using multiplicative weight updates based on the current allocation, and allocates to each agent by sampling from the exponential weights distribution as given by Lemma 5.3. The algorithm is parameterized by δ, the scale of the regularizer; by η, the rate at which the algorithm learns the dual variables $\boldsymbol{\alpha}$; and by scale parameter γ, which we set later.

Algorithm 4 Online Entropy Regularized Matching Algorithm (with parameters $\delta, \eta, \gamma \in \mathbb{R}_+$)

1: **input:** sample access to replica-surrogate matching instance values $\{v_{i,j}\}$ for replicas $i \in \{1, \ldots, mk\}$ and surrogates $j \in \{1, \ldots, m\}$.
2: **for all** $i \in \{1, \ldots, km\}$ **do**
3: Let k_j be the number of replicas previously matched to each surrogate j and $J = \{j : k_j < k\}$ the set of surrogates with availability remaining.
4: Set $\boldsymbol{\alpha}^{(i)}$ according to the exonential weights distribution with weights $\eta \cdot k_j$ for available surrogates $j \in J$ ($\alpha_j^{(i)} = 0$ for unavailable surrogates).
5: Match replica i to surrogate $j \in J$ drawn according to the exponential weights distribution with weights $(v_{i,j} - \gamma\, \alpha_j^{(i)})/\delta$ with the Fast Exponential Bernoulli Race (Algorithm 3).
6: **end for**

The algorithm needs to satisfy four properties to be useful in a polynomial time reduction. First, it needs to produce a perfect matching so that the replica-surrogate matching surrogate selection rule is stationary, specifically via condition (c). Second, it needs to be maximal-in-range for the real agent (replica i^*). In fact, all replicas are treated symmetrically and allocated by sampling from an exponential weights distribution that is maximal-in-range via Lemma 5.3. Third, it needs to have good welfare compared to the ideal matching. Fourth, its runtime needs to be polynomial. The first two properties are immediate and imply the theorem below. The last two properties are analyzed below.

Theorem 5.4. *The mechanism that maps types to surrogates via the replica-surrogate matching surrogate selection rule with the online entropy regularized matching algorithm (with payments from Theorem 4.1) is Bayesian incentive compatible.*

5.4 Social Welfare Loss

We analyze the welfare loss of the online entropy regularized matching algorithm (Algorithm 4) with regularizer parameter δ, learning rate η, and scale parameter γ set as a k-fraction of an estimate of the value of the offline program (Definition 5.3).

Theorem 5.5. *There are parameter settings for online entropy regularized matching algorithm (Algorithm 4) for which (1) its per-replica expected welfare is within an additive ϵ of the welfare of the optimal replica surrogate matching, and (2) given oracle access to $\mathcal{A}$, the running time of this algorithm is polynomial in m and $1/\epsilon$.*

To prove this theorem, we first argue how to set γ to be a constant approximation to the k-fraction of optimal value of the convex program with high probability, and with efficient sampling. Second, we argue that the online and offline optimal entropy regularized matching algorithms have nearly the same welfare. Finally, we argue that the offline optimal entropy regularized matching has nearly the welfare of the offline optimal matching. The proof of the theorem is then given by combining these results with the right parameters.

Parameter γ and approximating the offline optimal. Pre-setting γ to be an estimate of the optimal objective of the convex program in Definition 5.3 is necessary for the competitive ratio guarantee of Algorithm 4. Also, γ should be set in a symmetric and incentive compatible way across replicas, to preserve stationarity property. To this end, we look at an instance generated by an independent random draw of mk replicas (while fixing the surrogates). In such an instance, we estimate the expected values by sampling and taking the empirical mean for each edge in the replica-surrogate bipartite graph. We then solve the convex program exactly (which can be done in polynomial time using an efficient separation oracle). Obviously, this scheme is incentive compatible as we do not even use the reported type of true agent in our calculation for γ, and it is symmetric across replicas. In Appendix D we show how this approach leads to a constant approximation to the optimal value of the offline program in Definition 5.3 with high probability.

Lemma 5.6. *If $k = \Omega(\frac{\log(\eta^{-1})}{\delta^2 m(\log m)^2})$, then there exist a polynomial time approximation scheme to calculate γ (i.e. it only needs polynomial in m, k, δ^{-1} and η^{-1} samples to black-box allocation $\mathcal{A}$) such that*

$$\mathrm{OPT}(\mathbf{v})/k \le \gamma \le O(1)\,\mathrm{OPT}(\mathbf{v})/k$$

with probability at least $1-\eta$.

Competitive ratio of the online entropy regularized matching algorithm. Assuming γ is set to be a constant approximation to the k-fraction of the optimal value of the offline entropy regularized matching program, we prove the following lemma.

Lemma 5.7. *For a fixed regularizer parameter $\delta > 0$, learning rate $\eta > 0$, regularized welfare estimate γ, and market size $m \in \mathbb{N}$ that satisfy*

$$\frac{m \log m}{\eta^2} \le k \quad \text{and} \quad \mathrm{OPT}(\mathbf{v})/k \le \gamma \le O(1)\,\mathrm{OPT}(\mathbf{v})/k\ ,$$

the online entropy regularized matching algorithm (Algorithm 4) obtains at least an $(1 - O(\eta))$ fraction of the welfare of the optimal entropy regularized matching (Definition 5.3).

Proof. Recall that $\mathrm{OPT}(\mathbf{v})$ denotes the optimal objective value of the entropy regularized matching program. We will analyze the algorithm up to the iteration τ that the first surrogate becomes unavailable (because all k copies are matched to previous replicas).

Define the contribution of replica i to the Lagrangian objective of Observation 1 for allocation $\mathbf{x}_i = (x_{i,1}, \ldots, x_{i,m})$ and dual variables $\boldsymbol{\alpha}$ as

$$\mathcal{L}^{(i)}(\mathbf{x}_i, \boldsymbol{\alpha}) \triangleq \sum\nolimits_j v_{i,j}\, x_{i,j} - \delta \sum\nolimits_j x_{i,j} \log x_{i,j} + \sum\nolimits_j \gamma\alpha_j(\tfrac{1}{m} - x_{i,j}). \tag{3}$$

The difference between the outcome for replica i in the online algorithm and the solution to the offline optimization is that the algorithm selects the outcome with respect to dual variables $\gamma\boldsymbol{\alpha}^{(i)}$ while the offline algorithm selects the outcome with respect to the optimal dual variables $\boldsymbol{\alpha}^*$ (Observation 1). Denote the outcome of the online algorithm by

$$\mathbf{x}_i = (x_{i,1}, \ldots, x_{i,m}) = \mathrm{argmax}_{\mathbf{x}'_i \in \Delta^m}\, \mathcal{L}^{(i)}(\mathbf{x}'_i, \gamma\alpha^{(i)})\ ,$$

and its contribution to the objective by

$$\texttt{ALG}_i \triangleq \sum_j v_{i,j}\, x_{i,j} - \delta \sum_j x_{i,j} \log x_{i,j}\ .$$

Likewise for the outcome of the offline optimization by $\mathbf{x}_i^*$ and OPT_i. Denote by $\hat{\mathbf{x}}_i$ the indicator vector for the online algorithm sampling from $\mathbf{x}_i$.

Optimality of $\mathbf{x}_i$ for dual variables $\gamma\boldsymbol{\alpha}^{(i)}$ in equation (3) implies

$$\mathtt{ALG}_i + \sum\nolimits_j \gamma\alpha_j^{(i)}\left(\tfrac{1}{m} - x_{i,j}\right) \geq \mathrm{OPT}_i + \sum\nolimits_j \gamma\alpha_j^{(i)}\left(\tfrac{1}{m} - x^*_{i,j}\right)$$

so, by rearranging the terms and taking expectations conditioned on the observed history, we have

$$\begin{aligned}
\mathbf{E}[\mathtt{ALG}_i \mid \mathcal{H}_{i-1}] &\geq \gamma\mathbf{E}\Big[\boldsymbol{\alpha}^{(i)} \cdot \mathbf{x}_i \mid \mathcal{H}_{i-1}\Big] + \mathbf{E}[\mathrm{OPT}_i \mid \mathcal{H}_{i-1}] - \gamma\mathbf{E}\Big[\boldsymbol{\alpha}^{(i)} \cdot \mathbf{x}_i^* \mid \mathcal{H}_{i-1}\Big] \\
&= \mathbf{E}[\mathrm{OPT}_i] - \gamma\boldsymbol{\alpha}^{(i)} \cdot \mathbf{E}[\mathbf{x}_i^*] + \gamma\boldsymbol{\alpha}^{(i)} \cdot \hat{\mathbf{x}}_i - (\mathbf{E}[\mathrm{OPT}_i] - \mathbf{E}[\mathrm{OPT}_i \mid \mathcal{H}_{i-1}]) \\
&\quad + \gamma\boldsymbol{\alpha}^{(i)} \cdot (\mathbf{E}[\mathbf{x}_i^*] - \mathbf{E}[\mathbf{x}_i^* \mid \mathcal{H}_{i-1}]) + \gamma\boldsymbol{\alpha}^{(i)} \cdot (\mathbf{E}[\mathbf{x}_i \mid \mathcal{H}_{i-1}] - \hat{\mathbf{x}}_i) \\
&\geq \frac{1}{mk}\,\mathrm{OPT}(\mathbf{v}) + \gamma\boldsymbol{\alpha}^{(i)} \cdot (\hat{\mathbf{x}}_i - \frac{1}{m}\mathbf{1}) - L_i - L_i'
\end{aligned}$$

where

$$\begin{aligned}
L_i &\triangleq \gamma\ \boldsymbol{\alpha}^{(i)}.(\hat{\mathbf{x}}_i - \mathbf{E}[\mathbf{x}_i|\mathcal{H}_{i-1}])\ , \\
L_i' &\triangleq |(\mathbf{E}[r_i^*] - \mathbf{E}[r_i^*|\mathcal{H}_{i-1}]| + \gamma\|\mathbf{E}[\mathbf{x}_i^*] - \mathbf{E}[\mathbf{x}_i^*|\mathcal{H}_{i-1}]\|\ .
\end{aligned}$$

By summing the above inequalities for $i = 1 : \tau - 1$ we have:

$$\sum_{i=1}^{\tau-1} \mathbf{E}[\mathtt{ALG}_i|\mathcal{H}_{i-1}] \geq \frac{\tau-1}{mk}\,\mathrm{OPT}(\mathbf{v}) + \gamma\sum_{i=1}^{\tau-1}\boldsymbol{\alpha}^{(i)} \cdot (\hat{\mathbf{x}}_i - \frac{1}{m}\mathbf{1}) - \sum_{i=1}^{\tau-1}(L_i + L_i') \tag{4}$$

In order to bound the term $\gamma\sum_{i=1}^{\tau-1}\boldsymbol{\alpha}^{(i)} \cdot (\hat{\mathbf{x}}_i - \frac{1}{m}\mathbf{1})$, let $g_i(\boldsymbol{\alpha}) \triangleq \boldsymbol{\alpha} \cdot (\hat{\mathbf{x}}_i - \frac{1}{m}\mathbf{1})$. Then, by applying the regret bound of exponential gradient (or essentially multiplicative weight update) online learning algorithm for any realization of random variables $\{\hat{\mathbf{x}}_i\}$ (which will result in $\boldsymbol{\alpha}^{(i)}$ to be the exponential weights distributions with weights $\eta \cdot k_j$), we have

$$\sum_{i=1}^{\tau-1} g_i(\boldsymbol{\alpha}^{(i)}) \geq (1-\eta)\max_{\|\boldsymbol{\alpha}\|_1\leq 1, \boldsymbol{\alpha}\geq\mathbf{0}} \sum_{i=1}^{\tau-1} g_i(\boldsymbol{\alpha}) - \frac{\log m}{\eta} \geq (1-\eta)(k - \frac{\tau-1}{m}) - \frac{\log m}{\eta} \tag{5}$$

where the last inequality holds because at the time $\tau-1$, either there exists j such that $\sum_{i=1}^{\tau-1}\hat{x}_{i,j} = k$, or $\tau - 1 = mk$ and all surrogate outcome budgets are exhausted. In the former case, we have

$$\max_{\|\boldsymbol{\alpha}\|_1\leq 1, \boldsymbol{\alpha}\geq\mathbf{0}} \sum_{i=1}^{\tau-1} g_i(\boldsymbol{\alpha}) \geq \sum_{i=1}^{\tau-1} g_i(\mathbf{e}_j) \geq k - \frac{\tau-1}{m},$$

and in the latter case we have

$$\max_{\|\boldsymbol{\alpha}\|_1\leq 1, \boldsymbol{\alpha}\geq\mathbf{0}} \sum_{i=1}^{\tau-1} g_i(\boldsymbol{\alpha}) \geq 0 \geq k - \frac{\tau-1}{m}.$$

Combining (4) and (5), letting $Q_i = L_i + L_i'$, and assuming $\hat{\mathbf{x}}_j = 0$ for $j \geq \tau$, we have:

$$\begin{aligned}
\sum_{i=1}^{mk} \mathbf{E}[\mathtt{ALG}_i|\mathcal{H}_{i-1}] \geq \sum_{i=1}^{\tau-1} \mathbf{E}[\mathtt{ALG}_i|\mathcal{H}_{i-1}] &\geq \frac{\tau-1}{mk}\,\mathrm{OPT}(\mathbf{v}) + \gamma(1-\eta)(k - \frac{\tau-1}{m}) - \gamma\frac{\log m}{\eta} - \sum_{i=1}^{\tau-1} Q_i \\
&\geq \frac{\tau-1}{mk}\,\mathrm{OPT}(\mathbf{v}) + \frac{\mathrm{OPT}(\mathbf{v})}{k}(1-\eta)(k - \frac{\tau-1}{m}) - O(1)\cdot\mathrm{OPT}(\mathbf{v})\frac{\log m}{k\eta} - \sum_{i=1}^{mk} Q_i \\
&\geq (1-\eta)\,\mathrm{OPT}(\mathbf{v}) - O(\eta)\cdot\mathrm{OPT}(\mathbf{v}) - \sum_{i=1}^{mk} Q_i
\end{aligned} \tag{6}$$

where the last inequality holds simply because $k > \frac{\log m}{\eta^2}$. By taking expectations from both sides, we have

$$\mathbf{E}[\texttt{ALG}] \geq (1 - O(\eta)) \cdot \mathrm{OPT}(\mathbf{v}) - \sum_{i=1}^{mk} (\mathbf{E}[L_i] + \mathbf{E}\big[L_i'\big]) \tag{7}$$

We now bound each term separately. Define $Y_i \triangleq \sum_{i' \leq i} L_i$. Note that $\mathbf{E}[Y_i - Y_{i-1} | \mathcal{H}_{i-1}] = 0$, and therefore sequence $\{Y_i\}$ forms a martingale. Now, by using concentration of martingales we have the following lemma.

Lemma 5.8. $\mathbf{E}\Big[\sum_{i=1}^{mk} L_i\Big] \leq \gamma O(\sqrt{km \log km})$.

Proof of Lemma 5.8. Sequence $\{Y_i\}_{i=1}^{mk}$ forms a martingale, as $\mathbf{E}[Y_i - Y_{i-1} | \mathcal{H}_{i-1}] = 0$ and using Cauchy-Schwarz

$$|Y_i - Y_{i-1}| = \gamma |\boldsymbol{\alpha}^{(i)} \cdot (\hat{\mathbf{x}}_i - \mathbf{E}[\mathbf{x}_i | \mathcal{H}_{i-1}])| \leq \gamma \|\boldsymbol{\alpha}^{(i)}\| \cdot \|\hat{\mathbf{x}}_i - \mathbf{E}[\mathbf{x}_i | \mathcal{H}_{i-1}]\| \leq 2\gamma.$$

By using Azuma's inequality, we have

$$\Pr\{|Y_{mk}| \geq t\} \leq \exp\big(-\tfrac{t^2}{4km\gamma^2}\big).$$

Let $t = \gamma\sqrt{2km \log(km)}$, then $\Pr\{|Y_{mk}| \geq \gamma\sqrt{2km \log(km)}\} \leq \frac{1}{\sqrt{km}}$. Therefore,

$$\mathbf{E}\left[\sum_{i=1}^{mk} L_i\right] \leq \mathbf{E}[|Y_{mk}|] \leq \gamma\sqrt{2km \log(km)} + \tfrac{1}{\sqrt{km}} \cdot 2\gamma km = \gamma O(\sqrt{km \log km}). \qquad \square$$

To bound the second term, we use an argument based on Lemma 4.1 in Agrawal and Devanur (2015). In fact, we have the following lemma.

Lemma 5.9. $\mathbf{E}\Big[\sum_{i=1}^{mk} L_i'\Big] \leq \gamma O(\sqrt{k \log m})$.

Proof of Lemma 5.9. Using Lemma 4.1 in Agrawal and Devanur (2015) with $S = \{v \in \mathbb{R}^m : v \leq \frac{1}{m}\mathbf{1}\}$, we have $\mathbf{E}\Big[\sum_{i=1}^{mk} L_i'\Big] \leq \gamma O(\sqrt{skm \log m})$ where $s = \max_{v \in S} \max_{j \in [m]} v_j$. Obviously, $s = \frac{1}{m}$, which completes the proof. $\square$

Using Lemmas 5.9 and 5.8, combined with the facts that $\gamma \leq O(1) \cdot \frac{\mathrm{OPT}(\mathbf{v})}{k}$ and $k \geq \frac{m \log m}{\eta^2}$, we have $\mathbf{E}\Big[\sum_{i=1}^{mk} L_i + L_i'\Big] \leq O(\eta)\,\mathrm{OPT}(\mathbf{v})$. Together with (7), we conclude that $\mathbf{E}[\texttt{ALG}] \geq (1 - O(\eta))\,\mathrm{OPT}(\mathbf{v})$. This holds conditioned on $\frac{\mathrm{OPT}(\mathbf{v})}{k} \leq \gamma \leq O(1) \cdot \frac{\mathrm{OPT}(\mathbf{v})}{k}$. Moreover, γ is calculated by sampling such that this event happens with probability at least $(1 - \eta)$, which completes the proof. $\square$

Lemma 5.10. *With parameter $\delta \geq 0$ the welfare (average for the replicas) of the optimal entropy regularized matching is within an additive $\delta \log m$ of the welfare of the optimal matching.*

Proof. The entropy $-\sum_{i,j} x_{i,j} \log x_{i,j}$ is non-negative and maximized with $x_{i,j} = 1/m$. The maximum value of the entropy term is thus $\delta mk \log m$. The optimal objective value of the entropy regularized matching exceeds that of the optimal matching; thus, its welfare is within an additvie $\delta mk \log m$ of the optimal matching. The average welfare per replica in the entropy regularized matching (recall, there are mk replicas) is within $\delta \log m$ of the average welfare per replica on the optimal matching. $\square$

We conclude the section by combining Lemmas 5.6, 5.7, and 5.10 to prove the main theorem.

Proof of Theorem 5.5. Let $\delta = \frac{\epsilon}{3} \cdot \frac{1}{\log m}$ and $\eta = \frac{\epsilon}{3} \cdot \frac{1}{c}$, where c is a constant such that competitive ratio of Algorithm 4 is at least $1 - c \cdot \eta$ (Lemma 5.7). Moreover, let $k = \frac{m \log m}{\eta^2} = O(\frac{m \log m}{\epsilon^2})$, to satisfy the required condition in Lemma 5.7. The per-replica welfare of Algorithm 4 is within an additive $\delta \log m = \epsilon/3$ of its entropy regularized matching objective value, which in turn is a $1 - c \cdot \eta = 1 - \epsilon/3$ approximation to the per-replica optimal value of the entropy regularized matching due to Lemma 5.6 and 5.7. Following Lemma 5.10, the per-replica optimal value of the entropy regularized matching is within an additive $\delta \log m = \epsilon/3$ of the per-replica expected welfare of the optimal matching. As the per-replica welfare is bounded by 1, the per-replica welfare of the Algorithm 4 is within an additive $\epsilon/3 + \epsilon/3 + \epsilon/3 = \epsilon$ of the per-replica expected welfare of the optimal matching, as claimed. Finally, due to Lemma 5.6 and the fact that k is polynomial in m and $1/\epsilon$, the algorithm's running time is polynomial in m and $1/\epsilon$. □

5.5 Approximation Scheme for the Ideal Model and Black-box Reductions

We summarize the approximation scheme for the ideal model of Hartline et al. (2015), described and proved in previous sections, by the following theorem.

Theorem 5.11. *There is an BIC approximation scheme for the ideal m-to-m replica-surrogate reduction: Given any algorithm $\mathcal{A}$, market size m, and loss tollerance ϵ, there is a mechanism that (a) is Bayesian incentive compatible, (b) has per-agent no more than a polynomial in m and $1/\epsilon$ number of elementary operations and black-box calls to the algorithm $\mathcal{A}$, and (c) has per-agent expected welfare that is at least additive ϵ of that of the mechanism for the ideal m-to-m replica-surrogate reduction for $\mathcal{A}$.*

As an example application of Theorem 5.11, we apply it to one of the settings considered by Hartline et al. (2015). The result of this application of the theorem is a mechanism that has per-agent welfare within ϵ of that of an original non-monotone algorithm $\mathcal{A}$. The following definition and theorem can be found in Hartline et al. (2015).

Definition 5.4 (e.g., Hartline et al., 2015)**.** The *doubling dimension of a metric space* is the smallest constant Δ such that every bounded subset S can be partitioned into at most 2 subsets, each having diameter at most half of the diameter of S.

Theorem 5.12 (Hartline et al., 2015)**.** *For any agent with type space with doubling dimension $\Delta \geq 2$ and algorithm $\mathcal{A}$, if*

$$m \geq \frac{1}{2\epsilon^{\Delta+1}}$$

then the mechanism from the ideal model m-to-m replica-surrogate reduction has per-agent welfare that is at least an additive 2ϵ of the the algorithm's per-agent welfare.

We now have the following immediate corollary by combining Theorem 5.11 with Theorem 5.12.

Corollary 5.13. *For n agents, type spaces with doubling dimension bounded by Δ, and any algorithm $\mathcal{A}$, there is a mechanism that (a) is Bayesian incentive compatible, (b) has per-agent expected welfare that is at least an additive ϵ of the expected welfare of $\mathcal{A}$, and (c) has polynomial runtime in $\epsilon^{-\Delta-1}$ and n, given access to black-box oracle $\mathcal{A}$ and samples from the agents' type distributions.*

Acknowledgments

The authors would like to thank Nikhil Devanur, Shipra Agrawal and Pooya Jalaly for helpful discussions and comments on various parts of the paper.

References

Shipra Agrawal and Nikhil R Devanur. Fast algorithms for online stochastic convex programming. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1405–1424. SIAM, 2015.

Shipra Agrawal, Zizhuo Wang, and Yinyu Ye. A dynamic near-optimal algorithm for online linear programming. *arXiv preprint arXiv:0911.2974*, 2009.

Aaron Archer, Christos Papadimitriou, Kunal Talwar, and Éva Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. *Internet Mathematics*, 1(2): 129–150, 2004.

Moshe Babaioff, Robert Kleinberg, and Aleksandrs Slivkins. Multi-parameter mechanisms with implicit payment computation. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, pages 35–52, 2013.

Moshe Babaioff, Robert Kleinberg, and Aleksandrs Slivkins. Truthful mechanisms with implicit payment computation. *J. ACM*, 62(2):10:1–10:37, May 2015. ISSN 0004-5411.

Ashwinkumar Badanidiyuru, Robert Kleinberg, and Aleksandrs Slivkins. Bandits with knapsacks. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 207–216. IEEE, 2013.

Xiaohui Bei and Zhiyi Huang. Bayesian incentive compatibility via fractional assignments. In *Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms*, pages 720–733. SIAM, 2011.

Xiao Alison Chen and Zizhuo Wang. A near-optimal dynamic learning algorithm for online matching problems with concave returns. *arXiv preprint arXiv:1307.5934*, 2013.

Nikhil R Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 29–38. ACM, 2011.

Jason D Hartline and Brendan Lucier. Bayesian algorithmic mechanism design. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 301–310. ACM, 2010.

Jason D Hartline and Brendan Lucier. Non-optimal mechanism design. *The American Economic Review*, 105(10):3102–3124, 2015.

Jason D. Hartline, Robert Kleinberg, and Azarakhsh Malekian. Bayesian incentive compatibility via matchings. *SODA*, 2011.

Jason D. Hartline, Robert Kleinberg, and Azarakhsh Malekian. Bayesian incentive compatibility via matchings. *Games and Economic Behavior*, 92(C):401–429, 2015.

Zhiyi Huang and Sampath Kannan. The exponential mechanism for social welfare: Private, truthful, and nearly optimal. In *Proceedings of the 53rd IEEE Symposium on Foundations of Computer Science*, pages 140–149. IEEE, 2012.

Mark Huber. Optimal linear bernoulli factories for small mean problems. *CoRR*, abs/1507.00843, 2015.

MS Keane and George L O'Brien. A bernoulli factory. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 4(2):213–219, 1994.

Thomas Kesselheim, Andreas Tönnis, Klaus Radke, and Berthold Vöcking. Primal beats dual on online packing lps in the random-order model. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 303–312. ACM, 2014.

Krzysztof Łatuszyński. The bernoulli factory, its extensions and applications. *Proceedings of IWAP 2010*, pages 1–5, 2010.

Rajeev Motwani and Prabhakar Raghavan. *Randomized algorithms*. Chapman & Hall/CRC, 2010.

Şerban Nacu and Yuval Peres. Fast simulation of new coins from old. *The Annals of Applied Probability*, 15(1A):93–115, 2005.

A Implicit Payment Computation

In this section we describe one standard reduction for computing implicit payments in our general setting, given access to a BIC allocation algorithm $\tilde{\mathcal{A}}$: a multi-parameter counterpart of the single-parameter payment computation procedure used for example by Archer et al. (2004); Hartline and Lucier (2010), which makes $n+1$ calls to $\tilde{\mathcal{A}}$, thus incurring a factor $n+1$ overhead in running time. A different implicit payment computation procedure, described in Babaioff et al. (2013, 2015), avoids this overhead by calling $\tilde{\mathcal{A}}$ only once in expectation, but incurs a $1-\epsilon$ loss in expected welfare and potentially makes payments of magnitude $\Theta(1/\epsilon)$ *from* the mechanism *to* the agents.

The implicit payment computation procedure assumes that the agents' type spaces $(\mathcal{T}^k)_{k\in[n]}$ are *star-convex at 0*, meaning that for any agent k, any type $t^k \in \mathcal{T}^k$, and any scalar $\lambda \in [0,1]$, there is another type $\lambda t^k \in \mathcal{T}^k$ with the property that $v(\lambda t^k, o) = \lambda v(t^k, o)$ for every $o \in \mathcal{O}$. (The assumption is without loss of generality, as argued in the next paragraph.) The implicit payment computation procedure, applied to type profile $\mathbf{t}$, samples $\lambda \in [0,1]$ uniformly at random and computes outcomes $o^0 \triangleq \tilde{\mathcal{A}}(\mathbf{t})$ as well as $o^k \triangleq \tilde{\mathcal{A}}(\lambda t^k, \mathbf{t}^{-k})$ for all $k \in [n]$. The payment charged to agent k is $v(t^k, o^0) - v(t^k, o^k)$. Note that, in expectation, agent k pays

$$p^k(\mathbf{t}) = v(t^k, \tilde{\mathcal{A}}(\mathbf{t})) - \int_0^1 v(t^k, \tilde{\mathcal{A}}(\lambda t^k, \mathbf{t}^{-k}))\, d\lambda,$$

in accordance with the payment identity for multi-parameter BIC mechanisms when type spaces are star-convex at 0; see Babaioff et al. (2013) for a discussion of this payment identity.

Finally, let us justify the assumption that $\mathcal{T}^k$ is star-convex for all k. This assumption is without loss of generality for the allocation algorithms $\tilde{\mathcal{A}}$ that arise from the RSM reduction, because we can enlarge the type space $\mathcal{T}^k$ if necessary by adjoining types of the form λt^k with $t^k \in \mathcal{T}^k$ and $0 \le \lambda < 1$. Although the output of the original allocation algorithm $\mathcal{A}$ may be undefined when its

input type profile includes one of these artifically-adjoined types, the RSM reduction never inputs such a type into $\mathcal{A}$. It only calls $\mathcal{A}$ on profiles of surrogate types sampled from the type-profile distribution F, whose support excludes the artificially-adjoined types. Thus, even when the input to $\tilde{\mathcal{A}}$ includes an artifically-adjoined type λt^k, it occurs as one of the replicas in the reduction. The behavior of algorithm $\tilde{\mathcal{A}}$ remains well-defined in this case, because replicas are only used as inputs to the valuation function $v(r_i, o_j)$, whose output is well-defined even when $r_i = \lambda t^k$ for $\lambda < 1$.

B Basic Notions of Bayesian Mechanism Design

Multi-parameter Bayesian setting. Suppose there are n agents, where agent k has private *type* t^k from *type space* $\mathcal{T}^k$. The *type profile* of all agents is denoted by $\mathbf{t} = (t^1, \ldots, t^n) \in \mathcal{T}^1 \times \ldots \times \mathcal{T}^n$. Moreover, we assume types are drawn independently from known prior distributions. For agent k, let F^k be the distribution of $t^k \in \mathcal{T}^k$ and $\mathbf{F} = F^1 \times \ldots \times F^n$ be the joint distribution of types. Suppose there is an *outcome space* denoted by $\mathcal{O}$. Agent k with type t^k has valuation $v(t^k, o)$ for outcome $o \in \mathcal{O}$, where $v : (\mathcal{T}^1 \cup \ldots \cup \mathcal{T}^n) \times \mathcal{O} \to [0, 1]$ is a fixed function. Note that we assume agent values are non-negative and bounded, and w.l.o.g in $[0, 1]$. Finally, we allow charging agents with non-negative money *payments* and we assume agents are *quasi-linear*, i.e. an agent with private type t has *utility* $u = v(t, o) - p$ for the outcome-payment pair (o, p).

Algorithms, mechanisms and interim rules. An *allocation algorithm* $\mathcal{A}$ is a mapping from type profiles $\mathbf{t}$ to outcome space $\mathcal{O}$. A (direct revelation) mechanism $\mathcal{M}$ is a pair of *allocation rule* and *payment rule* $(\mathcal{A}, \mathbf{p})$, in which $\mathcal{A}$ is an allocation algorithm and $\mathbf{p} = (p^1, \ldots, p^n)$ where each p^k (denoted by the payment rule for agent k) is a mapping from type profiles $\mathbf{t}$ to $\mathbb{R}_+$. In fact, one can think of the interaction between strategic agents and a mechanism as following: agents submit their *reported types* $\mathbf{s} = (s^1, \ldots, s^n)$ and then the mechanism $\mathcal{M}$ picks the outcome $o = \mathcal{A}(\mathbf{s})$ and charges each agent k with its payment $p^k(\mathbf{s})$. We also consider *interim allocation rule*, which is the allocation from the perspective of one agent when the other agent's types are drawn from their prior distribution. More concretely, we abuse notation and define $\mathcal{A}^k(s_k) \triangleq \mathcal{A}(s^k, \mathbf{t}^{-k})$ to be the distribution over outcomes induced by $\mathcal{A}$ when agent k's type is s^k and other agent types are drawn from $\mathbf{F}^{-k}$. Similarly, for agent k we define *interim payment rule* $p^k(s^k) \triangleq \mathbf{E}_{\mathbf{t}^{-k} \sim \mathbf{F}^{-k}}\left[p^k(s^k, \mathbf{t}^{-k})\right]$, and *interim value* $v^k(s^k) \triangleq \mathbf{E}_{\mathbf{t}^{-k} \sim \mathbf{F}^{-k}}\left[v(s^k, \mathcal{A}^k(s^k, \mathbf{t}^{-k}))\right]$. In most parts of this paper, we focus only on one agent, e.g. agent k, and we just work with the interim allocation algorithm $\mathcal{A}^k(.)$. When it is clear from the context, we drop the agent's superscript, and therefore $\mathcal{A}(s)$ denotes the distribution over outcomes induced by $\mathcal{A}(s, \mathbf{t}^{-k})$ when $\mathbf{t}^{-k} \sim \mathbf{F}^{-k}$.

Bayesian and dominant strategy truthfulness. We are only interested in designing mechanisms that are *interim truthful*, i.e. every agent bests of by reporting her true type assuming all other agent's reported types are drawn independently from their prior type distribution. More precisely, a mechanism $\mathcal{M}$ is *Bayesian Incentive Compatible (BIC)* if for all agents k, and all types $s^k, t^k \in \mathcal{T}^k$,

$$\mathbf{E}_{\mathbf{t}^{-k} \sim \mathbf{F}-k}\Big[v(t^k, \mathcal{A}^k(t^k))\Big] - p^k(t^k) \geq \mathbf{E}_{\mathbf{t}^{-k} \sim \mathbf{F}-k}\Big[v(t^k, \mathcal{A}^k(s^k))\Big] - p^k(s^k) \tag{8}$$

As a stronger notion of truthfulness than Bayesian truthfulness, one can consider *dominant strategy truthfulness.*More precisely, a mechanism $\mathcal{M}$ is *Dominant Strategy Incentive Compatible (DSIC)* if for all agents k, and all types $s^k, t^k \in \mathcal{T}^k$ and all types $\mathbf{t}^{-k} \in \mathcal{T}^{-k}$,

$$v(t^k, \mathcal{A}(\mathbf{t})) - p^k(\mathbf{t}) \geq v(t^k, \mathcal{A}(s^k, \mathbf{t}^{-k})) - p^k(s^k, \mathbf{t}^{-k}) \tag{9}$$

Moreover, an allocation algorithm $\tilde{\mathcal{A}}$ is said to be BIC (or DSIC) if there exists a payment rule $\tilde{\mathbf{p}}$ such that $\tilde{M} = (\tilde{\mathcal{A}}, \tilde{\mathbf{p}})$ is a BIC (or DSIC) mechanism. Throughout the paper, we use the terms Bayesian (or dominant strategy) truthful and Bayesian (or dominant strategy) incentive compatible interchangeably. For randomized mechanisms, DSIC and BIC solution concepts are defined by considering expectation of utilities of agents over mechanism's internal randomness.

Social welfare. We are considering mechanism design for maximizing *social welfare*, i.e. the sum of the utilities of agents and the mechanism designer. For quasi-linear agents, this quantity is in fact sum of the valuations of the agents under the outcome picked by the mechanism. For the allocation algorithm $\mathcal{A}$, we use the notation $\text{val}(\mathcal{A})$ for the expected welfare of this allocation and $\text{val}^k(\mathcal{A})$ for the expected value of agent k under this allocation, i.e. $\text{val}(\mathcal{A}) \triangleq \mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\big[\sum_k v(t^k, \mathcal{A}(\mathbf{t}))\big]$ and $\text{val}^k(\mathcal{A}) \triangleq \mathbf{E}_{\mathbf{t}\sim\mathbf{F}}\big[v(t^k, \mathcal{A}(\mathbf{t}))\big]$.

C Surrogate Selection and BIC Reduction

Lemma C.1. *If matching algorithm $M(\mathbf{r}, \mathbf{s})$ produces a perfect k-to-1 matching for the instance in Definition 5.2, then its corresponding surrogate selection rule, denoted by Γ^M, is stationary*

Proof. Each surrogate s_j is an i.i.d. sample from F. Moreover, by the principle of deferred decisions the index i^* (the real agent's index in the replica type profile) is a uniform random index in $[mk]$, even after fixing the matching. Since this choice of replica is uniform in $[mk]$ and M is a perfect k-to-1 matching, the selection of surrogate outcome is uniform in $[m]$, and therefore the selection of surrogate type associated with this outcome is also uniform in $[m]$. As a result, the output distribution of the selected surrogate type is F. □

Lemma C.2. *If $M(\mathbf{r}, \mathbf{s})$ is a feasible replica-surrogate k-to-1 matching and is a truthful allocation rule (in expectation over allocation's random coins) for all replicas (i.e. assuming each replica is a rational agent, no replica has any incentive to misreport), then the composition of Γ^M and interim allocation algorithm $\mathcal{A}(.)$ forms a BIC allocation algorithm for the original mechanism design problem.*

Proof. Each replica-agent $i \in [mk]$ (*including* the real agent i^*) bests off by reporting her true replica type under some proper payments. Now, consider an agent in the original mechanism design problem with true type t. For any given surrogate type profile $\mathbf{s}$, using the Γ^M-reduction the agent receives the same outcome distribution as the one he gets matched to in M in a Bayesian sense, simply because of stationary property of Γ^M (Lemma C.1). As allocation M is incentive compatible, this agent doesn't benefit from miss-reporting her true type as long as the value he receives for reporting t' is $v(t, \mathcal{A}(\Gamma^M(t')))$. Therefore conditioning on $\mathbf{s}$ and non-real replicas in $\mathbf{r}$, the final allocation is BIC from the perspective of this agent. The lemma then follows by averaging over the random choice of $\mathbf{s}$ and non-real agent replicas in $\mathbf{r}$. □

D Estimating the Offline Optimal Regularized Matching

To formalize the approximation scheme, first fix the surrogate type profile $\mathbf{s}$. For a given replica profile $\mathbf{r}$ and replica-surrogate edge (i, j), let $v_{i,j}(r_i) = \mathbf{E}[v(r_i, \mathcal{A}(s_j))]$ and $\hat{v}_{i,j}(r_i)$ be the empirical mean of N samples of the random variable $v(r_i, \mathcal{A}(s_j)$. Suppose $\mathbf{v}(\mathbf{r})$ and $\hat{\mathbf{v}}(\mathbf{r})$ be the corresponding

vectors of expected values and empirical means under replica profile $\mathbf{r}$. Now, draw $\mathbf{r}'$ independently at random from the distribution of $\mathbf{r}$. We now show that $\text{OPT}(\hat{\mathbf{v}}(\mathbf{r}'))$ is a constant-factor approximation to $\text{OPT}(\mathbf{v}(\mathbf{r}))$ with high probability, and therefore we can use $\text{OPT}(\hat{\mathbf{v}}(\mathbf{r}'))$ to set γ.

We prove this in two steps. In Lemma D.1 we show for a given $\mathbf{r}'$, $\text{OPT}(\hat{\mathbf{v}}(\mathbf{r}'))$ is a constant-factor approximation to $\text{OPT}(\mathbf{v}(\mathbf{r}'))$ with high probability over the randomness in $\{\mathcal{A}(s_j)\}$. Then, in Lemma D.3 we show if $\mathbf{r}$ and $\mathbf{r}'$ are two random independent draws from the replica profile distribution then $\text{OPT}(\mathbf{v}(\mathbf{r}'))$ is a constant-factor approximation to $\text{OPT}(\mathbf{v}(\mathbf{r}))$ with high probability over randomness in $\mathbf{r}$ and $\mathbf{r}'$. These two pieces together prove our claim.

Lemma D.1. *If* $N \geq \frac{2\log(4m^2k\cdot\eta^{-1})}{\delta^2(\log m)^2}$, *then* $1/2\cdot\text{OPT}(\mathbf{v}(\mathbf{r}')) \leq \text{OPT}(\hat{\mathbf{v}}(\mathbf{r}')) \leq 2\,\text{OPT}(\mathbf{v}(\mathbf{r}'))$ *with probability at least* $1-\eta/2$.

Proof. By using the standard Chernoff-Hoeffding bound together with the union bound, with probability at least $1-2m^2ke^{-\frac{\delta^2(\log m)^2\cdot N}{2}} \geq 1-\eta/2$ we have

$$\forall(i,j)\in[km]\times[m]: \quad |\hat{v}_{i,j}(r'_i)-v_{i,j}(r'_i)| \leq 1/2\cdot\delta\log m$$

Suppose $\mathbf{x}^*$ is the optimal solution of the regularized matching convex program with values $\mathbf{v}(\mathbf{r}')$ and $\mathbf{x}^{**}$ is the optimal solution with values $\hat{\mathbf{v}}(\mathbf{r}')$.

$$\begin{aligned}\text{OPT}(\hat{\mathbf{v}}(\mathbf{r}')) &= \sum_i (\mathbf{x}_i^{**}\cdot\hat{\mathbf{v}}_i + \delta H(\mathbf{x}_i^{**})) \geq \sum_i (\mathbf{x}_i^{*}\cdot\hat{\mathbf{v}}_i + \delta H(\mathbf{x}_i^{*}))\\ &\geq \sum_i (\mathbf{x}_i^{*}\cdot\mathbf{v}_i + \delta H(\mathbf{x}_i^{*})) - \frac{\delta km\log m}{2} = \text{OPT}(\mathbf{v}(\mathbf{r}')) - \frac{\delta km\log m}{2} \geq 1/2\cdot\text{OPT}(\mathbf{v}(\mathbf{r}'))\end{aligned}$$

where the last inequality holds as $\text{OPT}(\mathbf{v}(\mathbf{r}'))$ is bounded below by the value of the uniform allocation, i.e. $\text{OPT}(\mathbf{v}(\mathbf{r}')) \geq \delta\cdot mk\log(m)$. Similarly, one can show $\text{OPT}(\mathbf{v}(\mathbf{r}')) \geq 1/2\cdot\text{OPT}(\hat{\mathbf{v}}(\mathbf{r}'))$, which completes the proof. □

Before proving the second step, we prove the following lemma, which basically shows that the optimal value of regularized matching $\text{OPT}(\mathbf{v}(\cdot))$ is a 1-Lipschitz multivariate function.

Lemma D.2. *For every* $i\in[km]$, *replica profile* $\mathbf{r}$, *and replica type* r'_i *we have:*

$$|\text{OPT}(\mathbf{v}(r_i,r_{-i})) - \text{OPT}(\mathbf{v}(r'_i,r_{-i}))| \leq 1$$

Proof. Let $\mathbf{x}$ and $\mathbf{x}'$ be the optimal assignments in $\text{OPT}(\mathbf{v}(r_i,r_{-i}))$ and $\text{OPT}(\mathbf{v}(r'_i,r_{-i}))$ respectively. We have

$$\begin{aligned}\text{OPT}(\mathbf{v}(r_i,r_{-i})) &= \sum_l(\mathbf{x}_l\cdot\mathbf{v}_l(r_l)+\delta H(\mathbf{x}_l)) \geq \sum_{l\neq i}(\mathbf{x}'_l\cdot\mathbf{v}_l(r_l)+\delta H(\mathbf{x}'_l)) + \mathbf{x}'_i\cdot\mathbf{v}_i(r_i)+\delta H(\mathbf{x}'_i)\\ &\geq \sum_{l\neq i}(\mathbf{x}'_l\cdot\mathbf{v}_l(r_l)+\delta H(\mathbf{x}'_l)) + \mathbf{x}'_i\cdot\mathbf{v}_i(r'_i)+\delta H(\mathbf{x}'_i) - 1 = \text{OPT}(\mathbf{v}(r'_i,r_{-i})) - 1\end{aligned}$$

where the last inequality holds because $\mathbf{x}'_i\cdot(\mathbf{v}_i(r'_i)-\mathbf{v}_i(r_i)) \leq 1$. Similarly, $\text{OPT}(\mathbf{v}(r'_i,r_{-i})) \geq \text{OPT}(\mathbf{v}(r_i,r_{-i})) - 1$ by switching the roles of r_i and r'_i. □

Lemma D.3. *If* $k \geq \frac{32\log(8\eta^{-1})}{\delta^2 m(\log m)^2}$, *then* $1/2\cdot\text{OPT}(\mathbf{v}(\mathbf{r})) \leq \text{OPT}(\mathbf{v}(\mathbf{r}')) \leq 3/2\cdot\text{OPT}(\mathbf{v}(\mathbf{r}))$ *with probability at least* $1-\eta/2$.

Proof. We start by defining the following Doob martingale sequence (Motwani and Raghavan, 2010), where (conditional) expectations are taken over the randomness in replica profile $\mathbf{r}$:

$$X_0 = \mathbf{E}[\mathrm{OPT}(\mathbf{v}(\mathbf{r}))]$$
$$X_n = \mathbf{E}[\mathrm{OPT}(\mathbf{v}(\mathbf{r}))|r_1, \cdots, r_n], \quad n = 1, 2, \ldots, km$$

It is easy to check that $\mathbf{E}[X_n|r_1, \ldots, r_{n-1}] = X_{n-1}$, and therefore $\{X_n\}$ forms a martingale sequence with respect to $\{r_n\}$. Moreover, $|X_n - X_{n-1}| \leq 1$ because of Lemma D.2. Now, by using Azuma–Hoeffding bound for martingales, we have

$$\Pr\{|X_{km} - X_0| \geq \delta km \log(m)/4\} \leq 2e^{\frac{-km\delta^2(\log m)^2}{32}}$$

and thus with probability at least $1 - 2e^{\frac{-km\delta^2(\log m)^2}{32}}$, $|\mathrm{OPT}(\mathbf{v}(\mathbf{r})) - \mathbf{E}[\mathrm{OPT}(\mathbf{v}(\mathbf{r}))]| \leq \frac{\delta km \log(m)}{4}$. Similarly, with probability at least $1 - 2e^{\frac{-km\delta^2(\log m)^2}{32}}$, we have $|\mathrm{OPT}(\mathbf{v}(\mathbf{r}')) - \mathbf{E}[\mathrm{OPT}(\mathbf{v}(\mathbf{r}'))]| \leq \frac{\delta km \log(m)}{4}$. Note that $\mathrm{OPT}(\mathbf{v}(\mathbf{r}))$ and $\mathrm{OPT}(\mathbf{v}(\mathbf{r}'))$ are identically distributed, and in particular they have the same expectation. Therefore with probability at least $1 - 4e^{\frac{-km\delta^2(\log m)^2}{32}}$ we have $|\mathrm{OPT}(\mathbf{v}(\mathbf{r})) - \mathrm{OPT}(\mathbf{v}(\mathbf{r}'))| \leq \frac{\delta km \log(m)}{2}$. By using the lower bound of $\delta km \log(m)$ for $\mathrm{OPT}(\mathbf{v}(\mathbf{r}))$ (due to uniform assignment), we conclude that with probability at least $1 - 4e^{\frac{-km\delta^2(\log m)^2}{32}} \geq 1 - \eta/2$ we have the following, as desired:

$$1/2 \cdot \mathrm{OPT}(\mathbf{v}(\mathbf{r})) \leq \mathrm{OPT}(\mathbf{v}(\mathbf{r}')) \leq 3/2 \cdot \mathrm{OPT}(\mathbf{v}(\mathbf{r})).$$ □

Corollary D.4. *If* $N \geq \frac{2\log(4m^2k\cdot\eta^{-1})}{\delta^2(\log m)^2}$ *and* $k \geq \frac{32\log(8\eta^{-1})}{\delta^2 m(\log m)^2}$*, then* $\gamma = \frac{4}{k}\mathrm{OPT}(\hat{\mathbf{v}}(\mathbf{r}'))$ *satisfies*

$$1/k \cdot \mathrm{OPT}(\mathbf{v}(\mathbf{r})) \leq \gamma \leq 12/k \cdot \mathrm{OPT}(\mathbf{v}(\mathbf{r})) \ ,$$

with probability at least $1 - \eta$*.*